arXiv:1210.4820v5 [cs.CR] 26 Nov 2012

Whole Genome Sequencing: Innovation Dream or Privacy Nightmare?

Emiliano De Cristofaro, *Palo Alto Research Center*
me@emilianodc.com

November 22, 2012

Abstract

Over the past several years, DNA sequencing has emerged as one of the driving forces in life-sciences, paving the way for affordable and accurate whole genome sequencing. As genomes represent the entirety of an organism's hereditary information, the availability of complete human genomes prompts a wide range of revolutionary applications. The hope for improving modern healthcare and better understanding the human genome propels many interesting and challenging research frontiers. Unfortunately, however, the proliferation of human genomes amplifies worrisome privacy concerns, since a genome represents a treasure trove of highly personal and sensitive information. In this article, we provide an overview of positive results and biomedical advances in the field, and discuss privacy issues associated with human genomic information. Finally, we survey available privacy-enhancing technologies and list a number of open research challenges.

1 Introduction

Over the past half a century, DNA sequencing has been one of the most active and fast-paced areas of research in life-sciences, yielding complete sequencing of many eukaryotic organisms, including men [1, 2]. A key, revolutionary role in this context has been played by **High-Throughput Sequencing (HTS)** techniques. In 2007, scientists sequenced the first diploid human genome in 2007 [3] and recently completed a project to sequence 1,000 human genomes [4].

The $3B, 13-year Human Genome Project [2] has involved a number of research institutions worldwide and is considered one of the major breakthroughs of this century. Nowadays, different HTS technologies are competing to accurately sequence an individual human genome, composed of about **3 billion DNA nucleotides**, with prices affordable for a large number of individuals.

The race for **cheaper and more accurate whole genome sequencing technologies** has been quite exciting, plunging costs from $1B only a decade ago to $250,000 in 2008 (by Illumina), and to about $4,400 a couple of years ago (in 2009 by Complete Genomics [5] and in 2011, again, by Illumina [6]). Life Technologies announced this year that they can scan the full genome for $1,000 [7], while Oxford Nanopore announced their intent to commercialize a sequencer the size of a USB memory stick that can sequence a whole genome for $900 in 15 minutes [8]. Geniachip claims to go beyond this and deliver the same results for just $100. Large corporations, such as IBM and GE have also entered the race in the last couple of years, while, recently, Roche has unsuccessfully attempted to acquire Illumina.

Nowadays, the landscape of companies and technologies competing in this sector is so fast-paced that it becomes relatively hard to keep a comprehensive, up-to-date list. Nonetheless, it is evident that whole genome sequencing will be a reality in the near future, a commodity costing less than an X-ray or an MRI scan.

In this article, we discuss how advances in whole genome sequencing have created **contrasting feelings** pertaining the implications of widespread availability of whole genomes. On the one hand, the hope for improving modern healthcare and better understanding the human genome has attracted significant research investments and, arguably, generated many groundbreaking results. On the other hand, however, a number of alarming privacy and ethical concerns has been raised

pertaining to the sensitivity of human genomic information and its disclosure.

We provide an overview of positive results and recent biomedical advances in the field (Sec. 2), and discuss privacy issues associated with human genomic information (Sec. 3). Finally, we provide a list of a few compelling research challenges in the area (Sec. 4) and survey the state of the art in privacy-enhancing technologies focusing on computational genomic tests (Sec. 5).

2 The good news: Beyond Personalized Medicine

Undoubtedly, ubiquity of human genomes creates enormous opportunities and challenges. In particular, it promises to launch a new era of genome-enabled predictive, preventive, participatory, and personalized medicine ("P4 medicine") [9].

Personalized Medicine is recognized as a significant paradigm shift and a major trend in health care, moving us closer to a more precise, powerful, and holistic type of medicine [10]. With personalized medicine, treatment and medication type/dosage would be tailored to the precise genetic makeup of individual patients. Experts predict that advances in whole genome sequencing will further stimulate development of personalized medicine [11]. Commercial companies like Knome already offer services that take raw genome data and create usable reports for doctors. In general, the availability of a patient's fully sequenced genome will enable clinicians, doctors, and testing facilities to run a number of complex, correlated genetic tests in a matter of seconds, using specialized computational algorithms (as opposed to more expensive and slower in-vitro tests).

Already today, personalized medicine is a reality in a number of medical scenarios. Measurements of *erbB2* protein in breast, lung, or colorectal cancer patients are taken before selecting proper treatment. It has been shown that the trastuzumab monoclonal antibody is effective only in patients whose genetic receptor is over-expressed [12]. Also, testing for the thiopurine S-methyltransferase (*tpmt*) gene is required prior to prescribing for 6-mercaptopurine and azathioprine – two drugs used for treating childhood leukemia and autoimmune diseases. The *tpmt* gene codes for the TPMT enzyme that metabolizes thiopurine drugs: genetic polymorphisms affecting enzymatic activity are correlated with variations in sensitivity and toxicity response to such drugs. Patients suffering from this genetic disease (1 in 300) only need 6-10% of the standard dose of thiopurine drugs; if treated with the full dose, they risk severe bone marrow suppression and subsequent death [13]. Similarly, doctors who want to prescribe Zelboraf (Roche's treatment for advanced skin cancer) first test the patient for the BRAFV 600E mutation, which is found in about half of all cases. Other analogous examples include the Philadelphia chromosome mutations related to Acute Lymphoblastic Leukemia (ALL) and BRCA1/BRCA2 genes in correlation to familial breast and ovarian cancer syndromes.

Experts estimate that about a third of the **900 cancer drugs** currently in clinical trials could soon come to market with a DNA or other molecular test attached [14]. Although most predominant, cancer treatment is only one of the application fields of personalized medicine. For instance, a recent Canadian study has shown how, for some cardiac patients, recovery from a common heart procedure can be complicated by a single gene responsible for drug processing, and that selection of blood thinner drugs should depend on whether or not patient holds such a gene mutation [15]. Also, in a study that shows how whole-genome sequencing could be used in life-or-death medical situations involving newborns, researchers at a hospital in Kansas analyzed the entire genomes of seven babies that died near birth, accurately diagnosing five of them with critical conditions within about 50 hours each – fast enough to be meaningful to their care [16].

Tremendous advances in **Pharmacogenomics** (the study of the impact of genetic variation on the response to medications) are also driving research in the field. Examples include genes involved in the action and metabolism of warfarin (coumadin), a medication used as an anticoagulant [17], as well as genes encoding Cytochrome P450 enzymes, which metabolize neuroleptic medications to improve drug response and reduce side-effects [18].

The availability of whole human genomes will also facilitate a number of genetic tests that today are performed in vitro, such as, paternity/ancestry testing and genetic compatibility, by reducing costs and time.

For instance, **genetic paternity tests** may be run very efficiently in computation, by designing algorithms that emulate in-vitro, highly accurate, court-admissible tests, e.g., based on Short Tandem Re-

peats (STRs) and Restriction Fragment Length Polymorphisms (RFLPs).[1] Actually, running those algorithms on whole genomes could even improve the accuracy of paternity tests: experts point out that, while any two unrelated humans share about 99.5% of their genomes, individuals tied by a parent-child relationship have 99.8% of their genome in common. Thus, one could realize error-free paternity tests by counting the number of matching nucleotides across test takers.

On a similar note, **ancestry and genealogical testing** allows individuals to trace their lineage by analyzing their genomic information (the scope of such tests being often quite heterogeneous.) Ancestry testing is useful in a myriad of health-related applications (e.g., susceptibility to diseases common to certain populations) but it is also increasingly used in social or recreational scenarios, e.g., to map one's own genetic heritage or find common ancestry. Several commercial entities (e.g., 23andMe [19]) already maintain a collection of sample genomes from individuals belonging to different ethnic groups, and compare them against their customers' genomic information to understand how they relate to known ethnic groups.

Genetic compatibility tests let potential or existing partners assess the possibility of transmitting to their children a genetic disease with Mendelian inheritance [20]. Modern genetic testing can accurately predict whether a couple is at risk of conceiving a child with an autosomal recessive disease. Consider, for instance, *Beta-Thalassemia minor*, that causes red cells to be smaller than average, due to a mutation in the *hbb* gene. It is called *minor* when the mutation occurs only in one allele. This *minor* form has no severe impact on a subject's quality of life. However, the *major* variant—that occurs when both alleles carry the mutation—is likely to result in premature death, usually, before age twenty. Therefore, if both partners silently carry the *minor* form, there is a 25% chance that their child could carry the major variety.

In general, **genetic tests** are routinely used for several purposes, such as newborn screening, confirmational diagnostics, as well as pre-symptomatic testing, e.g., predicting Huntington's disease [21] and estimating risks of various congenital diseases. In fact, 23andMe [19] provides relatively low-cost genetic tests for 960,000 specific Single-Nucleotide Polymorphisms (SNPs). (SNPs are the most common form of DNA variation occurring when a single nucleotide differs between members of the same species or paired chromosomes of an individual [22, 23].) However, while some diseases (e.g., Huntington's) are caused by mutations in a single gene and are easily tested *in vitro*, the risk of developing other diseases depends on multiple genes, which makes them difficult to identify. Low-cost genetic sequencing provides researchers with much more genomic information, and enables them to identify new genetic variations as well as run more complicated tests.

While the relationship between advances in whole genome sequencing and breakthroughs in personalized medicine and genetic tests creates a lot of research "enthusiasm", a number of biomedical experts have also expressed doubts related to the limits of gene mapping's power to predict a person's likelihood of developing a disease [24]. It remains unclear how the availability of large numbers of whole genomes will yield a better understanding of the human genome (and correlated diseases), e.g., through **Genome-Wide Association Studies (GWAS)**. These studies examine common genetic variants in a very large set of individuals to find out if any variant is associated with, e.g., a disease, and possibly correlate the disease to a given ancestry line. Additional areas to explore include genetic compatibility tests for sperm and organ donors [25], evolutionary studies (e.g., based on genomes of Denisovans and Neanderthals [26]), as well as research on genomes of crops and animals [27].

3 The bad news: Privacy, Legal, and Ethical Concerns

Widespread and low-cost availability of HTS technologies and genomic data has raised a number of ethical, security, and privacy concerns [28]. The human genome not only uniquely and irrevocably identifies its owner, but also contains information about ethnic heritage, predisposition to numerous diseases and conditions, including mental disorders, and many other phenotypic traits [29, 30, 31]. Recent studies suggest that even political preferences may be influenced by voters' genetic makeup [32]. Furthermore, due to its hereditary nature, disclosing one's human genome also implies, to a certain extent, disclosing the genomes of close relatives.

[1] E.g., in RFLP-based paternity test, individuals' genomes are probed and cut by enzyme digestion and the test outcome is assessed based on the similarity of resulting fragments.

Traditional approaches to privacy, such as de-identification or aggregation [33], become completely moot in the genomic era, since the genome itself is the **ultimate identifier**. To further compound the privacy problem, health information is increasingly shared electronically among insurance companies, health care providers and employers. This, coupled with the possibility of creating large centralized genome repositories (e.g., for GWAS research), raises the specter of possible abuses. (A few results exploring de-anonymizing individuals from genomic datasets include [34, 35].)

Long before whole genome sequencing prices drop to a few hundred dollars, society had already envisioned a future where the issue of genetic discrimination could dramatically affect social dynamics, hiring and healthcare practices, and even ways of procreating. Even popular culture, with sci-fi movies and narrative literature, has expressed its concerns – for instance, the concept of **genism** actually originated from the 1997 Hollywood movie "Gattaca" [36], denoting the theory that distinctive human characteristics and abilities are determined by genes, based on DNA sequence characteristics with resulting in discrimination as pernicious as racism [37]. (One could note how genism actually shares several traits with eugenic ideals prominent in the hateful policies of some regimes, e.g., the Third Reich.) The movie led molecular biologist Lee M. Silver to write in Nature Genetics that "Gattaca is a film that all geneticists should see if for no other reason than to understand the perception of our trade held by so many of the public-at-large" [38].

Several funding agencies, e.g., the US National Human Genome Research Institute (NHGRI), has established—from the very beginning of the Human Genome Project, in 1990—efforts like the **Ethical, Legal and Social Implications** (ELSI) Research Program, to foster basic and applied research on the ethical, legal and social implications of genetic and genomic research for individuals, families and communities. Some **federal laws** have been passed to start addressing privacy issues. The 2003 Health Insurance Portability and Accountability Act (HIPAA) provides a general framework for protecting and sharing Protected Health Information (PHI), and, in 2008, the Genetic Information Nondiscrimination Act (GINA) was adopted to prohibit discrimination on the basis of genetic information with respect to health insurance and employment [39].

While providing general guidelines and a basic safety net, current legislation does not offer detailed technical information about safe and privacy-preserving ways for storing and querying genomes. Privacy practitioners are strongly advocating the **need for more restrictive legislation** as a result of gaps in current policies – see, e.g., a comprehensive list of EPIC's efforts at http://epic.org/privacy/genetic/. Also, a very recent report from the Presidential Commission for the Study of Bioethical Issues [40] has analyzed advances of whole genome sequencing, and highlighted growing concerns about privacy and security. The report lists 12 privacy and security recommendations, including de-identification. (On a separate note, while the report is related to our article, observe that its scope is quite different from ours: we aim at a technical analysis of technologies and challenges, whereas, [40] provides a high-level "effort to identify and promote policies and practices that ensure scientific research, health care delivery, and technological innovation are conducted in a socially and ethically responsible manner".)

At the policy level, the main open challenges include, for instance, the need for **informed consent** to guard against surreptitious DNA testing by requiring authorities and companies to obtain written permission from citizens before collecting, analyzing, storing or sharing their genetic information (e.g., preventing people from collecting hair or saliva samples and maliciously sequencing the victims' genome).

On the other hand, some academic researchers fear that privacy-restrictive measures could seriously hinder genomic research. Scientists typically sequence DNA from thousands of people to discover genes associated with particular diseases, thus, the informed consent restriction would mean that large genomic datasets could not be re-used to study a different disease – researchers would either need to destroy the data after each study, or track down thousands of former subjects for new authorizations [41].

4 Research Challenges

While privacy issues are not yet hampering progress in basic genomic research, it is not too early to investigate them, as discussed above, in light of their complexity and potential impact on society.

In order for computational genetic tests on whole human genomes to become accepted and commonplace,

efficient and (possibly) privacy-preserving versions of such tests need to be developed. This poses a number of challenges, which we investigate below.

Accessibility: As we discussed in Section 1, it is reasonable to assume that, in a few years, a relevant number of individuals worldwide will have access to their fully-sequenced genome. Due to its sensitivity and size (about 3 billion letters), one of the most difficult research challenges is related to how and where should the genomes be stored. Should it be given to the individual and stored on her PC? On a USB stick? On dedicated hardware? On her smartphone? Or should the genome be trusted with another entity? A physician? A healthcare provider? The health insurance provider? A trusted third-party cloud? Naturally, answering these questions requires a clear understanding of information technology as well as legal, ethical, privacy, and ethnographic issues (which are closely connected to challenges discussed below).

Privacy: Given its extreme sensitivity, an individual should, ideally, never disclose personal genomic information. However, one should be able to allow others (e.g., individuals, doctors, clinicians) to run specific genetic tests that yield nothing beyond their intended results. For instance, letting a testing facility run some genetic tests should rather not result in surrendering one's whole genome.

In this context, additional motivations (besides ethical and legal ones) for privacy protection stem from liability concerns. Mere possession of a patient's sensitive information would require the testing entity to demonstrate that the information was treated appropriately and disposed of when no longer needed. Considering several recent (and rather frequent) incidents of massive losses of sensitive records, the entity might be unwilling to assume additional risk.

Long-term data safety: The human genome uniquely identifies its owner, but also discloses a lot of information about its relatives as well as its descendants, even several generations into the future. This prompts the problem of long-term data safety, even if human genomes are always stored encrypted. An encryption scheme considered strong today might gradually weaken in the long term. Consequently, it is not too far-fetched to imagine that a third-party in possession of an encrypted genome might be able to decrypt it, e.g., 20 or 50 years later. Whereas, genome sensitivity does not dissipate over time.

Accuracy and Accountability: Computational genomic tests should guarantee accuracy and reliability comparable to current (and widely accepted) lab-based *in-vitro* equivalents. For example, a software implementation of the paternity test on fully sequenced genomes should offer at least the same confidence as its *in-vitro* counterpart, currently admissible in a court of law. Also, computational tests should aim at accountability, e.g., by providing guarantees that tests are run correctly and on intended genomic information.

Efficiency: Computational genomic tests should incur minimal storage, communication, and computational costs. Arguably, minimality in this setting is relative to the context of such tests. For instance, patients may be inclined (and used) to wait several days to obtain results of genetic tests that concern their health, however, in the computational setting, long waiting times might hinder the real-world practicality of these tests (besides taking out one of the main motivations for computational tests.) Also, if a patient's genome is stored on her PC or phone, usability of these tests will be minimized, e.g., due to connectivity and battery life issues.

Usability: Computational genomic tests that involve end-users should be usable by, and meaningful to, regular non-tech-savvy individuals. This translates into non-trivial questions, such as: how much understanding should be expected from a user running a test? What information (and at what level of granularity) should be presented to the user as part of a test and as its outcome? Do privacy perceptions and concerns experienced by patients correspond to what the scientific community would expect? For instance, one may think that patients will be likely to trade off privacy of their genome to enable tests that can save them from, e.g., cancer. However, to the best of our knowledge, no scientific study has investigated users' concerns and (mis)proven common beliefs to this regard, thus, pointing out the need for ethnographic studies in the field. Also, it remains an open problem to explore effective ways to communicate to the users the potential privacy risks associated with genomic information and its disclosure.

Large-scale research on human genomes: Finally, as discussed in Section 3, potential privacy, legal, and ethnical concerns appear contrasting to enabling large-scale research on human genomes, such as, Genome-

Wide Association Studies (GWAS). One of the necessary conditions for effective GWAS is actually the large availability of human genomes, e.g., in order to discover correlations between genetic makeups and medical conditions. Consequently, a current research trend is to store donors' genomes on clouds and clusters of computers and employ *"big data"* mining and searching technologies (such as, MapReduce) for genomic research initiatives, like GWAS, sequencing alignment, etc. [42, 43]. Once again, related privacy and legal concerns (also related to relatives and descendants) remain a challenging open problem.

5 Available Techniques for Privacy-Preserving Tests on Genomic Data

Motivated by the extreme sensitivity of genomic information, the security research community has been attuned to the emergence of whole genome sequencing and a few privacy-preserving cryptographic techniques have been proposed in recent literature. Alas, the majority of them focus on (and only apply to) secure computation on DNA fragments, and not to whole genomes. However, a couple of recent results have begun investigating privacy-respecting tests on whole genomes.

5.1 Secure Computation on DNA Fragments

Results on privacy-preserving operations on DNA fragments are mainly of two kinds: (1) secure searching/matching DNA strings, (2) computing the similarity of DNA sequences. We review them in the following.

Troncoso-Pastoriza, et al. [44] proposed an error-resilient privacy-preserving protocol for string searches. One party, on input of a DNA snippet, can verify the existence of a short template (e.g., a genetic test held by the service provider) within its (short) snippet. This technique handles errors and maintains privacy of both the template and the snippet. Each query is represented as an automaton executed using a finite state machine (FSM) in an oblivious manner. Also, secure pattern matching techniques, e.g., those in [45] and [46], could also be applied to securely search binary strings in a DNA snippet. Then, Katz, et al. [47] realized secure computation of the CODIS test [48] (run by the FBI for DNA identity testing) and other search tests that could not be otherwise implemented using pattern matching or FSM. Alas, the communication and computational complexities introduced by cryptographic operations in these techniques are not practical for real-world deployment (even worse if one considers applying these techniques to whole genomes).

Another set of cryptographic results focus on privately computing the *edit distance* for two strings, or DNA snippets, α, β. Recall that edit distance is defined as the minimum number of operations, such as, delete, insert, or replace, needed to transform α into β. Privacy-preserving computation of Smith-Waterman scores [49] has also been investigated and used for sequence alignment. Jha, et al. [50] (and follow-up work) show how to securely compute edit distance using garbled circuits [51], but demonstrate that the resulting overhead is acceptable only for small strings.

Wang, et al. [52] developed techniques for computation on genomic data stored at a data provider, including: edit distance, Smith-Waterman and search for homologous genes. Program specialization is used to partition genomic data into "public" (most of the genome) and "sensitive" (a very small subset of the genome). Sensitive regions are replaced with symbols by data providers before data consumers have access to genomic information. However, due to today's limited understanding of human genomes, such partition in sensitive and non-sensitive information is likely to be completely ineffective in a few years.

Finally, Franz et al. [53] show how privacy of genomic sequences can be protected while they are analyzed using Hidden Markov Models (HMM). HMM is commonly used in bioinformatics to detect certain non-beneficial patterns in the genome and, unsurprisingly, allows more powerful computations than string matching.

5.2 Secure Testing on Whole Human Genomes

Baldi, et al. [54] recently introduced several cryptographic protocols for privacy-preserving testing of whole human genomes, including paternity tests and genetic screening for personalized medicine or recessive genetic diseases. In their setting, individuals obtain their genomes and allow authorized parties (e.g., doctors and clinicians) to run genetic tests such that only test results are disclosed to one or both parties (with provable security). However, [54] only addresses the issue of designing cryptographic pro-

tocols, and does not deal with issues like where to store whole sequences. To this end, their follow-up work [55] starts tackling such challenges, and proposes a framework and an implemented toolkit, called *GenoDroid*. It incorporates several techniques offering efficient privacy-preserving genomic testing and runs on commodity Android smartphones – it is available at http://sprout.ics.uci.edu/projects/privacy-dna. Also, preliminary user studies seem to support usability and acceptability of proposed techniques.

Recent work in [56, 57] also explore privacy-enhancing technologies for medical tests on genomic data and build a secure architecture for privacy-preserving disease-susceptibility tests.

Finally, Chen, et al. [42] studied the problem of privacy-preserving mapping and aligning of human genomic sequences to a reference genome, by outsourcing work to a hybrid cloud. In fact, at sequencing time, human genomes are read in short sequences, and these need to be aligned by comparing them to a reference genome. The work in [42] enables one to perform this task by outsourcing the computation to cloud and protecting DNA information marked as sensitive.

6 Conclusions

This article presented an analysis of recent progress in whole genome sequencing. We first provided an overview of new technologies, applications, and biomedical advances, stimulated by the promise of widespread availability of complete human genomes. In particular, the hope for personalized medicine, i.e., tailoring diagnosis and treatment to patients' genetic makeup, has prompted a number of pioneering results. Then, we investigated privacy issues associated with human genomic information, as human genomes represent a treasure trove of highly personal and sensitive information. Finally, we surveyed the state of the art in privacy-enhancing technologies focusing on computational genomic tests and provided a compelling list of several research challenges that call for extensive work in this area.

Acknowledgments: The author would like to thank Marshall Bern, Stefan Katzenbeisser, and Jean-Pierre Hubaux for the useful discussions and valuable feedback. Also, the idea to write this article is inspired by the work published in [54] and [55], co-authored with Pierre Baldi, Roberta Baronio, Sky Faber, Paolo Gasti, and Gene Tsudik, to whom the author is immensely grateful – without their work, none of this would be possible.

References

[1] R. Waterston *et al.*, "Initial sequencing and comparative analysis of the mouse genome," *Nature*, vol. 420, no. 6915, 2002.

[2] International Human Genome Sequencing Consortium, "Initial sequencing and analysis of the human genome," *Nature*, vol. 409, 2001.

[3] S. Levy *et al.*, "The diploid genome sequence of an individual human," *PLoS biology*, vol. 5, no. 10, 2007.

[4] D. Altshuler *et al.*, "A map of human genome variation from population scale sequencing," *Nature*, vol. 467, no. 7319, 2010.

[5] R. Drmanac *et al.*, "Human genome sequencing using unchained base reads on self-assembling DNA nanoarrays," *Science*, vol. 327, no. 5961, 2010.

[6] K. Darcè, "Illumina drops sequencing price to $4,000." http://is.gd/illumina_drops_prices, 2011.

[7] S. Moore, "Moore's Law Leads to $1000 Genome Device," *IEEE Spectrum*, January 2012.

[8] K. Hallam, "Oxford Nanopore To Sell Tiny DNA Sequencer." http://bloom.bg/Aogf12, 2012.

[9] L. Hood and D. Galas, "P4 Medicine: Personalized, Predictive, Preventive, Participatory A Change of View that Changes Everything." http://www.cra.org/ccc/docs/init/P4_Medicine.pdf, 2009.

[10] A. Weston and L. Hood, "Systems biology, proteomics, and the future of health care: toward predictive, preventative, and personalized medicine," *Journal of Proteome Research*, vol. 3, no. 2, 2004.

[11] G. Ginsburg and H. Willard, "Genomic and personalized medicine: foundations and applications," *Translational Research*, vol. 154, no. 6, 2009.

[12] A. Prat and J. Baselga, "The role of hormonal therapy in the management of hormonal-receptor-positive breast cancer with co-expression of her2," *Nature Clinical Practice Oncology*, vol. 5, no. 9, 2008.

[13] A. Abbott, "Special section on human genetics: With your genes? Take one of these, three times a day," *Nature*, vol. 425, no. 6960, 2003.

[14] A. Burke, "Foundation Medicine: Personalizing Cancer Drugs." http://is.gd/foundation_medicine, 2012.

[15] S. Wood, "RAPID GENE: Point-of-care genetic test singles out clopidogrel nonresponders." http://www.theheart.org/article/1379471.do, 2012.

[16] K. Weintraub, "Sick Babies Could Have Genomes Sequenced in Days." http://is.gd/sick_babies_genomes_in_days, 2012.

[17] J. Oldenburg, M. Watzka, S. Rost, and C. Müller, "VKORC1: Molecular target of coumarins," *Journal*

of Thrombosis and Haemostasis, vol. 5, no. s1, 2007.
[18] S. Cichon, M. Nöthen, M. Rietschel, and P. Propping, "Pharmacogenetics of schizophrenia," *American Journal of Medical Genetics*, vol. 97, no. 1, 2001.
[19] 23andMe. https://www.23andme.com/.
[20] V. McKusick and S. Antonarakis, *Mendelian inheritance in man: a catalog of human genes and genetic disorders*. John Hopkins University Press, 1994.
[21] J. Gusella *et al.*, "A polymorphic DNA marker genetically linked to Huntington's disease," *Nature*, vol. 306, no. 5940, 1983.
[22] P. Stenson *et al.*, "The human gene mutation database: 2008 update," *Genome Medicine*, vol. 1, no. 1, 2009.
[23] National Center for Biotechnology Information (US), "Single Nucleotide Polymorphism Database." http://www.ncbi.nlm.nih.gov/projects/SNP/.
[24] G. Naik, "Gene Maps Are No Cure-All." http://on.wsj.com/HackSD, 2012.
[25] J. Gitlin, "Using DNA Sequencing to Monitor Organ Transplant Rejection." http://www.genome.gov/27544325, 2012.
[26] M. Meyer *et al.*, "A high-coverage genome sequence from an archaic denisovan individual," *Science*, 2012.
[27] J. Beckmann and M. Soller, "Restriction fragment length polymorphisms and genetic improvement of agricultural species," *Euphytica*, vol. 35, no. 1, 1986.
[28] F. Collins and V. McKusick, "Implications of the Human Genome Project for medical science," *Jama*, vol. 285, no. 5, 2001.
[29] J. Fowler, J. Settle, and N. Christakis, "Correlated genotypes in friendship networks," *Proceedings of the National Academy of Sciences*, vol. 108, no. 5, 2011.
[30] T. Canli, "The emergence of genomic psychology," *Nature*, vol. 8, 2007.
[31] M. Fumagalli *et al.*, "Parasites represent a major selective force for interleukin genes and shape the genetic predisposition to autoimmune conditions," *Experimental Medicine*, vol. 206, no. 6, 2009.
[32] D. Benjamin *et al.*, "The genetic architecture of economic and political preferences," *Proceedings of the National Academy of Sciences*, vol. 109, no. 21, 2012.
[33] B. Malin, "An evaluation of the current state of genomic data privacy protection technology and a roadmap for the future," *Journal of the American Medical Informatics Association*, vol. 12, no. 1, 2005.
[34] R. Wang *et al.*, "Learning your identity and disease from research papers: information leaks in Genome Wide Association Study," in *CCS*, 2009.
[35] N. Homer *et al.*, "Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays," *PLoS Genetics*, vol. 4, no. 8, 2008.
[36] A. Niccol, "GATTACA." http://imdb.to/bRWLfh, 2007.
[37] G. Annas, "Genism, Racism, and the prospect of genetic genocide," in *World Conference Against Racism, Racial Discrimination, Xenophobia and Related Intolerance*, 2001.
[38] L. Silver, "Genetics goes to Hollywood," *Nature Genetics*, vol. 17, no. 3, 1997.
[39] M. Wadman, "Genetics bill cruises through senate," *Nature*, vol. 453, 2008.
[40] Presidential Commission for the Study of Bioethical Issues, "PRIVACY and PROGRESS in Whole Genome Sequencing." http://www.bioethics.gov/cms/sites/default/files/PrivacyProgress508.pdf, 2012.
[41] H. Shen, "California considers DNA privacy law – Academic researchers fear measures would prohibit work with genetic databases." *Nature*, http://www.nature.com/news/california-considers-dna-privacy-law-1.10677, 2012.
[42] Y. Chen, B. Peng, X. Wang, and H. Tang, "Large-Scale Privacy-Preserving Mapping of Human Genomic Sequences on Hybrid Clouds," in *NDSS*, 2012.
[43] R. Metz, "A Startup Uses the Cloud to Unravel DNA." http://is.gd/cloud_unraveling_dna, 2012.
[44] J. Troncoso-Pastoriza, S. Katzenbeisser, and M. Celik, "Privacy preserving error resilient dna searching through oblivious automata," in *CCS*, 2007.
[45] R. Gennaro, C. Hazay, and J. Sorensen, "Text Search Protocols with Simulation Based Security," in *PKC*, 2010.
[46] C. Hazay and T. Toft, "Computationally secure pattern matching in the presence of malicious adversaries," *Asiacrypt*, 2010.
[47] J. Katz and J. Malka, "Secure text processing with applications to private dna matching," in *CCS*, 2010.
[48] The Federal Bureau of Investigation, "Combined DNA Index System (CODIS)." http://www.fbi.gov/about-us/lab/codis, 2011.
[49] T. Smith and M. Waterman, "Identification of common molecular subsequences," *Journal of Molecular Biology*, vol. 147, 1981.
[50] S. Jha, L. Kruger, and V. Shmatikov, "Towards practical privacy for genomic computation," in *S&P*, 2008.
[51] A. Yao, "Protocols for secure computations," in *FOCS*, 1982.
[52] R. Wang, X. Wang, Z. Li, H. Tang, M. Reiter, and Z. Dong, "Privacy-preserving genomic computation through program specialization," in *CCS*, 2009.
[53] M. Franz, B. Deiseroth, K. Hamacher, S. Jha, S. Katzenbeisser, and H. Schröder, "Towards secure bioinformatics services (short paper)," *Financial Cryptography and Data Security*, 2012.
[54] P. Baldi, R. Baronio, E. De Cristofaro, P. Gasti, and G. Tsudik, "Countering GATTACA: Efficient and Secure Testing of Fully-Sequenced Human Genomes," in *CCS*, 2011.
[55] E. De Cristofaro, S. Faber, P. Gasti, and G. Tsudik, "GenoDroid: Are Privacy-Preserving Genomic Tests Ready for Prime Time?," in *WPES*, 2012.

[56] E. Ayday, J. L. Raisaro, and J.-P. Hubaux, "Privacy-Enhancing Technologies for Medical Tests Using Genomic Data." https://documents.epfl.ch/users/a/ay/ayday/www/erman_publications/genomic_privacy_paper.pdf, 2012.

[57] E. Ayday, M. Humbert, J. Fellay, P. J. McLaren, J. Rougemont, J. L. Raisaro, A. Telenti, and J.-P. Hubaux., "Privacy-Enhancing Technologies for Medical Tests Using Genomic Data." https://documents.epfl.ch/users/a/ay/ayday/www/erman_publications/genomic_privacy.pdf, 2012.